

04-17-00

A

1 IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

jc79  
04/14/00  
2 Inventorship.....Hendel et al.  
3 Applicant.....Microsoft Corporation  
4 Attorney's Docket No. ....MS1-468US  
5 Title: Methods and Arrangements for Generating Debugging Information Following Software Failures

6 **TRANSMITTAL LETTER AND CERTIFICATE OF MAILING**

7 To: Commissioner of Patents and Trademarks,  
8 Washington, D.C. 20231

9 From: Thomas A. Jolly (Tel. 509-324-9256; Fax 509-323-8979)  
10 Lee & Hayes, PLLC  
11 421 W. Riverside Avenue, Suite 500  
12 Spokane, WA 99201

13 The following enumerated items accompany this transmittal letter and are being submitted for the  
14 matter identified in the above caption.

15 1. Specification-title page, plus 35 pages, including 66 claims and Abstract  
16 2. Transmittal letter including Certificate of Express Mailing  
17 3. 6 Sheets Formal Drawings (Figs. 1-6)  
18 4. Return Post Card

19 Large Entity Status [x] Small Entity Status [ ]

20 Date: 4-14-2000

21 By: Thomas A. Jolly  
22 Thomas A. Jolly  
23 Reg. No. 39,241

24 **CERTIFICATE OF MAILING**

25 I hereby certify that the items listed above as enclosed are being deposited with the U.S. Postal  
Service as either first class mail, or Express Mail if the blank for Express Mail No. is completed below, in  
an envelope addressed to The Commissioner of Patents and Trademarks, Washington, D.C. 20231, on the  
below-indicated date. Any Express Mail No. has also been marked on the listed items.

Express Mail No. (if applicable) EL580803869

Date: 4.14.00

By: Lori A. Vierra  
Lori A. Vierra

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Methods and Arrangements for Generating Debugging  
Information Following Software Failures**

Inventor:

Matthew Hendel

&

Kent Forschmiedt

ATTORNEY'S DOCKET NO. MS1-468US

1

## **TECHNICAL FIELD**

2 This invention relates to computers and software, and more particularly to  
3 methods and arrangements for efficiently generating debugging information  
4 following a software failure.

5

## **BACKGROUND OF THE INVENTION**

6 Computer failures can generally be traced to either hardware or software  
7 problems. This Background section discusses those failures that can be identified  
8 through a careful examination of the data stored by the computer at the time of  
9 failure.

10 When an operating system or application program crashes it is useful, if  
11 possible, to save information about the reasons for the crash and the state of the  
12 system when it crashed for later reference. Conventional techniques for collecting  
13 such POSTMORTEM INFORMATION (or DUMP INFORMATION) require an  
14 enormous amount of data be stored. For example, when an operating system  
15 (OS) crashes, the common technique of collecting postmortem information is to  
16 save the entire contents of the computer's RAM to permanent storage (e.g., disk,  
17 tape, floppy, etc.).

18 As computer's memory sizes and the amount of data associated with the OS  
19 continues to increase, the time it takes to store this postmortem information upon  
20 failure is correspondingly increased. Indeed, many users simply don't have time  
21 to generate postmortem information, and instead opt to manually reboot the  
22 computer, losing all postmortem debugging data. Consequently, problems may or  
23 may not be reported, and those that are reported would lack the critical data  
24 necessary to debug the problem. This means that a random, yet commonly

1 occurring problem may not get the attention required. For example, consider the  
2 effort required to generate a complete postmortem debug information file for a  
3 large file server with 64 gigabytes of system memory. It would likely take the  
4 system many hours to store all of this postmortem information, if enough disk  
5 space could even be found to store it on. And if data was finally stored to  
6 memory, it would very difficult to move later. Even with very fast networks, it  
7 would require hours to copy a 64GB file over a network connection. Even a  
8 conventional personal computer (PC) having only 64 megabytes of system  
9 memory could still take an inordinate amount of time to complete a full  
10 postmortem dump.

11 To avoid such delays, some operating systems are configured to output only  
12 that portion of the system memory that is allocated for use by the operating system  
13 kernel. While this tends to reduce the amount of data generated, even the resulting  
14 postmortem dump file is quite large. For example, when configured to save only  
15 the kernel portion of the system memory, the postmortem files for the Microsoft  
16 Windows 2000 kernel range in size from 32 megabytes to 8 gigabytes in size. For  
17 reference, a 32-megabyte file would take about 3 hours to transfer over a 28.8K  
18 Baud modem connection.

19 This same problem occurs with non-operating system programs, also called  
20 user-mode programs. As with OS components, the main problem with user-mode  
21 post-mortem debug data is that it is typically quite large and it takes a long time to  
22 generate. User-mode dump files for Windows 2000 are typically 50 to 100  
23 megabytes in size. As we discussed above, with files this large is it very difficult  
24 to transmit files of this size back to the computer or operating system vendor for  
25 analysis.

1           Consequently, the above stated problems and conventional solutions  
2 hamper the desire of many users and manufacturers for improved online support of  
3 the OS and applications. Here, for example, it would be unacceptable and  
4 potentially expensive for a user having a 28.8K Baud modem to transmit a 64 MB  
5 memory dump file to the manufacturer for postmortem analysis (it would take  
6 more than 5 hours).

7           As such, there is a need for improved methods and arrangements that  
8 substantially reduce the requisite amount of data required to conduct a significant  
9 postmortem analysis following an operating system or application failure.  
10 Preferably, the methods and arrangements will be advantageously configured to  
11 allow for online user support for a variety of users, computing devices, operating  
12 systems, applications, and the like.

13 **SUMMARY OF THE INVENTION**

14           Improved methods and arrangements are provided that substantially reduce  
15 the requisite amount of data required to conduct postmortem analysis following an  
16 operating system or application failure. The methods and arrangements can be  
17 advantageously configured to allow for rapid online user support for a variety of  
18 users, computing devices, operating systems, applications, and the like. The  
19 methods can be applied to either operating system failures or application program  
20 failures.

21           The various methods and arrangements address the generation of failure  
22 information when failures occur at either the operating system level (KERNEL-  
23 MODE) or in traditional user program (USER-MODE). For example, when a  
24 failure occurs to a program executing in USER-MODE a USER-MODE

1 MINIDUMP or USER MINIDUMP can be generated. Similarly, when a failure  
2 occurs in kernel-mode a KERNEL-MODE MINIDUMP or KERNEL MINIDUMP  
3 can be generated.

4 The above stated needs and others are met by a method that includes  
5 determining when to generate a dump file, and generating a dump file by gathering  
6 the thread, thread context and callstack for the thread that caused the failure; also  
7 included in the dump are the process containing the failing thread and the reason  
8 for the crash (e.g., invalid memory reference). The resulting dump file can then be  
9 stored to a storage medium and accessed for subsequent analysis. In certain  
10 implementations the dump file also includes information about multiple threads,  
11 multiple processes or portions of data for the failing process. When configured as  
12 a kernel minidump, the list of device drivers for the system is included.

13 To further support a kernel minidump, the method may also include  
14 allocating a buffer space in memory during an initialization process, and reserving  
15 space on a storage medium drive suitable for writing the contents of the buffer.  
16 The method then includes generating the kernel minidump file by initially storing  
17 the list of device drivers, crashing process and thread, as well as the thread-context  
18 for the thread and the callstack associated with the thread. Additionally, the kernel  
19 minidump contains information identifying the reason for generating the crash.  
20 The method further includes copying the kernel memory dump file from the buffer  
21 space to the storage medium as a kernel minidump. In still other implementations,  
22 the method includes, upon re-initialization, accessing the kernel minidump on the  
23 storage medium and using at least a portion of it to further understand the failure  
24 that occurred.

1        The method is also applicable to non-operating system programs. Here, for  
2 example, a user minidump file is generated by gathering the following information  
3 at the time of a failure: the list of threads running at the time of the crash; the  
4 thread-context and callstack for each running thread; the list of modules loaded at  
5 the time of the crash; the reason for the crash; and, selected regions of the  
6 process's memory that relate to the cause of the crash.

7        The above stated needs and others are met by an arrangement having a  
8 system coupled to memory and at least one data storage device. The operating  
9 system is configured to determine when to generate a dump file while running  
10 computer instructions, and generate the dump file in the memory by gathering  
11 callstack information for at least one running thread, thread context information  
12 about the at least one running thread, process identifying information associated  
13 with the at least one running thread, and information identifying the reason for  
14 generating the dump file.

15       Methods for use in communicating between a client process and a server  
16 process in a distributed processing system are also provided. Here, for example,  
17 the methods can be used to provide application-programming interfaces (APIs) or  
18 like capability when writing and reading dump files.

19       For a write dump operation, such methods include having the client process  
20 issue a write dump file call having a plurality of call parameters including a  
21 process handle, a process identifier, a handle to a file where dump file information  
22 is to be written, and a dump type identifier. The server process receives this write  
23 dump file call and parses the call to retrieve the parameters, and then in response,  
24 issues a write dump file call acknowledgment providing a true-false indication.

1 For a read dump operation, such methods include having the client process  
2 issue a read dump file call having a plurality of call parameters including a header  
3 of a dump file and a data type identifier of data to read from a dump file, having  
4 the server process receiving the read dump file call and parsing the call to retrieve  
5 the parameters, and then having the server process issue a read dump file call  
6 acknowledgment providing a true-false indication and a plurality of call return  
7 parameters including a pointer to a beginning of a dump stream, and a stream size  
8 identifying the size of the dump stream. The plurality of call return parameters  
9 may also include a pointer to a dump file directory.

10

11 **BRIEF DESCRIPTION OF THE DRAWINGS**

12 A more complete understanding of the various methods and arrangements  
13 of the present invention may be had by reference to the following detailed  
14 description when taken in conjunction with the accompanying drawings wherein:

15 Fig. 1 is a block diagram depicting an exemplary computer system suitable  
16 for use in generating a dump file.

17 Fig. 2 is a flow-chart depicting a kernel minidump process for use in  
18 generating postmortem debug information of an operating system failure as might  
19 be experienced from time to time in the computer system of Fig. 1.

20 Fig. 3 is a flow-chart depicting a user minidump process for use in  
21 generating user postmortem debug information of a computer application failure,  
22 for example, as might be experienced from time to time in the computer system of  
23 Fig. 1.

1 Fig. 4 is a table depicting experimental results for user minidump files  
2 associated with different applications, as generated, for example, using the user  
3 minidump process of Fig. 3.

4 Fig. 5 is an illustrative representation of an exemplary user-mode  
5 minidump file having a plurality of directories, for example, as generated using the  
6 user-mode minidump process of Fig. 3.

7 Fig 6 is a block diagram depicting features of an exemplary operating  
8 system suitable for use in a user-mode minidump process or a kernel-mode  
9 minidump process.

10  
11 **DETAILED DESCRIPTION**

12 As shown in Fig. 1, computer 20 includes one or more processors or  
13 processing units 21, a system memory 22, and a bus 23 that couples various  
14 system components including the system memory 22 to processors 21. Bus 23  
15 represents one or more of any of several types of bus structures, including a  
16 memory bus or memory controller, a peripheral bus, an accelerated graphics port,  
17 and a processor or local bus using any of a variety of bus architectures.

18 The system memory includes read only memory (ROM) 24 and random  
19 access memory (RAM) 25. A basic input/output system (BIOS) 26, containing the  
20 basic routines that help to transfer information between elements within computer  
21 20, such as during start-up, is stored in ROM 24.

22 Computer 20 further includes a hard disk drive 27 for reading from and  
23 writing to a hard disk, not shown, a magnetic disk drive 28 for reading from and  
24 writing to a removable magnetic disk 29, and an optical disk drive 30 for reading  
25 from or writing to a removable optical disk 31 such as a CD ROM, DVD ROM or

1 other optical media. The hard disk drive 27, magnetic disk drive 28 and optical  
2 disk drive 30 are each connected to bus 23 by applicable interfaces 32, 33 and 34,  
3 respectively.

4 The drives and their associated computer-readable media provide  
5 nonvolatile storage of computer readable instructions, data structures, program  
6 modules and other data for computer 20. Although the exemplary environment  
7 described herein employs a hard disk, a removable magnetic disk 29 and a  
8 removable optical disk 31, it should be appreciated by those skilled in the art that  
9 other types of computer readable media which can store data that is accessible by a  
10 computer, such as magnetic cassettes, flash memory cards, digital video disks,  
11 random access memories (RAMs) read only memories (ROM), and the like, may  
12 also be used in the exemplary operating environment.

13 A number of program modules may be stored on the hard disk, magnetic  
14 disk 29, optical disk 31, ROM 24, or RAM 25, including an operating system 35,  
15 one or more application programs 36, other program modules 37, and program  
16 data 38. A user may enter commands and information into computer 20 through  
17 input devices such as keyboard 40 and pointing device 42. Other input devices  
18 (not shown) may include a microphone, joystick, game pad, satellite dish, scanner,  
19 or the like. These and other input devices are connected to the processing unit 21  
20 through an interface 46 that is coupled to bus 23.

21 A monitor 47 or other type of display device is also connected to bus 23 via  
22 an interface, such as a video adapter 48. In addition to the monitor, personal  
23 computers typically include other peripheral output devices (not shown) such as  
24 speakers and printers.

1 Computer 20 can operate in a networked environment using logical  
2 connections to one or more remote computers, such as a remote computer 50.  
3 Remote computer 50 may be another personal computer, a server, a router, a  
4 network PC, a peer device or other common network node, and typically includes  
5 many or all of the elements described above relative to computer 20. The logical  
6 connections depicted in Fig. 2 include a local area network (LAN) 51 and a wide  
7 area network (WAN) 52. Such networking environments are commonplace in  
8 offices, enterprise-wide computer networks, intranets, and the Internet.

9 When used in a LAN networking environment, computer 20 is connected to  
10 the local network 51 through a network interface or adapter 156. When used in a  
11 WAN networking environment, computer 20 typically includes a modem 54 or  
12 other means for establishing communications over the wide area network 52, such  
13 as the Internet. Modem 54, which may be internal or external, is connected to bus  
14 23 via interface 46. In a networked environment, program modules depicted  
15 relative to the personal computer 20, or portions thereof, may be stored in the  
16 remote memory storage device. It will be appreciated that the network  
17 connections shown are exemplary and other means of establishing a  
18 communications link between the computers may be used.

19 Figure 6 depicts a simplified model of an Operating System. At the top of  
20 the diagram is the operating system's KERNEL. The kernel runs it's own code  
21 and the code from any number of KERNEL DEVICE DRIVERS. Note that the  
22 arrow connecting the kernel to the device drivers is labeled "1 to N". This means  
23 that there can be many device drivers for every kernel. The kernel runs any  
24 number of PROCESSES. Each process contains a kernel portion and user portion.  
25 The kernel portion of a process is accessible only by the kernel and it's device

1 drivers. The user portion is to both the user-mode process and kernel-mode  
2 portion of the process.

3 The process's code is executed either by the kernel (KERNEL CODE) or  
4 by the process itself (USER CODE). Processes can accomplish several tasks at  
5 once. Each of these tasks is a THREAD. A thread, like a process, contains both a  
6 user-mode portion, and a private kernel-mode portion. Each thread contains one  
7 THREAD CONTEXT. A thread context specifies exactly what a thread is doing  
8 at any specific moment. The thread context contains one CALLSTACK. Thus,  
9 there is a kernel-mode only callstack and a user-mode callstack.

10 A flow-chart depicting an exemplary kernel-mode minidump process 100  
11 for use in generating data suitable for conducting a postmortem analysis of a  
12 computer failure is provided in Fig. 2.

13 As discussed below, kernel-mode minidump process 100 generates a very  
14 compact representation of the running system information in the event of a system  
15 failure. The resulting kernel minidump file is capable of being created and written  
16 to disk very quickly, as well as quickly transported, for example, via e-mail over  
17 the Internet or other communication resource. Despite its relatively small size, the  
18 kernel minidump file contains enough information to allow a developer or product  
19 support engineer to significantly reproduce the computing environment and  
20 otherwise debug the faulting driver or operating system failure.

21 One advantage of kernel minidump process 100 is that it is capable of  
22 directly addressing the reliability and stability of the operating system, since all  
23 system failures can be logged and sent to the manufacturer for analysis. Thus, for  
24 example, as a result of kernel minidump process 100, both device driver  
25 developers and operating system developers can easily diagnose and quickly

1 respond to user issues that usually only occur in the field. Furthermore, since the  
2 resulting kernel minidump file is significantly compact enough it can also be  
3 added to a database of problems requiring further analysis or comparison.

4 In addition to traditional operating systems, it should be noted that kernel  
5 minidump process 100 is also a particularly desirable feature for enterprise-level  
6 operating systems, especially those operating systems that use a large number of  
7 third party drivers or have high reliability requirements.

8 Unlike previous solutions, kernel minidump process 100 is preferably  
9 designed to generate as minimal amount of information as necessary to  
10 meaningfully debug the failure. Thus, in accordance with certain exemplary  
11 implementations, kernel minidump process 100 generates a kernel minidump file  
12 essentially containing the processor information for the crashing processor, the  
13 thread for the crashing process, the process information for the crashing process,  
14 the callstack for the crashing thread, a list of device drivers loaded at the time of  
15 the crash including the location within the system where they are loaded and the  
16 version of the device that was loaded, and the specific cause for the crash.

17 In this example, kernel minidump process 100 does not store the actual  
18 code (i.e., executable instructions) for the running operating system or any of the  
19 loaded drivers. Instead, kernel minidump process 100 stores version information  
20 about the operating system and each driver that was loaded, and the virtual address  
21 where the driver was loaded. Then, during the subsequent analysis/debugging  
22 stage, the analyst can locate the proper driver version and load it into memory for  
23 debugging purposes at that time.

24 By not saving code, therefore, kernel minidump process 100 is able to  
25 generate a much more compact kernel minidump file. By way of example, in

1 certain exemplary implementations the kernel minidump file for Microsoft®  
2 Windows® 2000 is only 64 KB, still contains sufficient information to locate the  
3 crashing driver, generate a symbolic callstack for the driver and locate the source  
4 code line and file where the driver failed.

5 In accordance with still further features, a user may also access certain local  
6 variables and information about the cause of the crash. These may be helpful to  
7 online or telephone support personnel or processes.

8 In step 102 of Fig. 2, at initialization/boot time, a kernel minidump buffer  
9 capable of holding the required kernel minidump information is allocated along  
10 with a kernel minidump file that is large enough to hold the kernel minidump  
11 information. For example, in the Windows® 2000 example, the kernel minidump  
12 file is allocated from the paging file. Also in step 102, a list of raw disk sectors  
13 that form the kernel minidump file is obtained.

14 Next, in step 104 space for the disk drivers is reserved on the disk drive  
15 where the kernel minidump file will be stored.

16 In step 106, when the computer system fails, the processor context of the  
17 crashing processor is gathered from the system and written to the kernel minidump  
18 buffer. In step 108, the thread information for the thread that was executing when  
19 the failure occurred is gathered from the system and written to the kernel  
20 minidump buffer. Similarly, in step 110, the process information for the process  
21 that owns the crashing thread (above) is gathered from the system and written to  
22 the kernel minidump buffer.

23 In step 112, the kernel stack for the crashing thread is gathered from the  
24 system and written to the kernel minidump buffer.

1        In Step 114 the list of kernel-mode components, such as, the kernel, device  
2        drivers and hardware abstraction layers (HALs) that were loaded when the  
3        computer failure occurred and their respective load-addresses and appropriate  
4        version information (e.g., a checksum, timestamp, and image size) is gathered  
5        from the system and written to the kernel minidump buffer.

6        In step 116, specific information about the cause of the crash is gathered  
7        from the system and also written to the kernel minidump buffer. In Windows®  
8        2000, for example, this information is provided as the BugCheckCode along with  
9        the four BugCheck words associated with the failure.

10       In step 118, the kernel minidump information from steps 106 through 116  
11       is transferred or otherwise copied from the kernel minidump buffer to disk. Since  
12       the file system is never invoked, faults in the file-system can be caught. In certain  
13       implementations, the resulting kernel minidump information is stored in an  
14       indexed table.

15       Next, in step 120, when the computer system is rebooted, the kernel  
16       minidump file is transferred from it's temporary location, where it was written to  
17       by fault-handling code in step 118, to a more permanent location, for example, in a  
18       user-specified directory containing a list of kernel minidump files. These files can  
19       then be accessed, forwarded or otherwise handled as necessary to study the failure.

20       Those skilled in the art will recognize that the various steps in kernel  
21       minidump process 100 can be rearranged or reconfigured in many ways, as  
22       required to meet the needs of a given computing system environment.

23       One of the interesting aspects about the resulting kernel minidump file is  
24       what information can be left out. Thus, in the example above, no code for running  
25       components is saved (i.e., executable instructions), no dynamically allocated

1 memory is saved, and no initialized or uninitialized data is saved. Also, no  
2 information about threads, processes or processors except the crashing thread,  
3 process and processor is saved in the example above. While certain information  
4 may be necessary for other implementations, in the above example (Fig. 2) this  
5 appears to significantly minimize the amount of information required.

6 Kernel minidump process 100 is essentially directed towards system  
7 crashes. Thus, there remains a need for similar processes for non-system failures,  
8 i.e., user application failures.

9 Conventional postmortem debug information collection techniques tend to  
10 store all user-accessible memory for an application, including code for all user-  
11 mode shared libraries, dynamically allocated memory, and all stack data, upon the  
12 occurrence of a failure. This produces very large dump files, of which only a  
13 portion may be required to subsequently analyze the failure.

14 With this in mind, Fig. 3 depicts a user minidump process 200 for use in  
15 generating significantly smaller user-mode dump files containing selected  
16 information that is suitable for conducting postmortem analysis of an application  
17 failure.

18 In step 202, user minidump process 200 gathers information about the  
19 callstack for all running threads. In step 204, the exception information that  
20 caused the fault, if any, and thread context of the fault is gathered. A list of all  
21 loaded modules for the faulting application is gathered in step 206. In step 208,  
22 the processor context for all threads is gathered. In some configurations of  
23 minidump process 200, additional information about the crash is also saved at this  
24 point. In step 210 the results of steps 202 through 208 are written to disk.  
25

1        In accordance with certain implementations, for example, user minidump  
2 process 200 does not require instrumentation of the application, hooking of system  
3 application programming interface (API) calls, or either local or remote symbolic  
4 information to be available. Instead, the resulting user minidump file can be  
5 generated by the client application simply using only information readily available  
6 through standard debugging APIs, such as, for example, Win32 APIs.

7        The current implementation of the Dr. Watson tool in Windows® 2000, for  
8 example, saves data for the entire process's address space. While this tends to  
9 guarantee full postmortem analysis of the application, huge dump files are often  
10 created (e.g., 50MB or more)

11        In contrast, using user minidump process 200, in the event that Microsoft®  
12 Internet Explorer (IE) crashes, instead of a Dr. Watson dialog coming up, IE can  
13 itself catch the error and write a user minidump file to disk. Since the user  
14 minidump file size is relatively small (about 60 KB for IE), it takes only a second  
15 or two to write the user minidump file to disk. When IE restarts, it could then note  
16 that a user minidump file was written. Here, for example, IE may then access a  
17 remote bug database that identifies that the problem experienced was a bug that  
18 has since been fixed in a Service Pack release of IE. The user could then be  
19 prompted to upgrade IE to fix the problem. In another illustrative scenario, assume  
20 Microsoft® Excel crashes on a user's machine. As with the IE scenario, with user  
21 minidump process 200 Excel can catch the exception and write a user minidump  
22 file to disk. Here, the user may be notified that there was a program error and that  
23 the user minidump file should be mailed to the developer. The minidump file can  
24 then be e-mailed to developer where it can be forwarded to the development team  
25 responsible for the code in Excel. Using standard debugging tools like Visual

1 Studio, the development team is able to walk the callstack, view parameters to  
2 functions, view local variables, and view the exact line and file where the failure  
3 occurred. Using this information the development team will likely be able to fix  
4 an obscure program error that otherwise would not have been reproducible.

5 In certain implementations of Windows® 2000, for example, two additional  
6 functions in the form of APIs are added to the DBGHELP system DLL. The first  
7 API is a write user minidump API (*WriteMiniDump*), and the second API is a read  
8 user minidump API (*ReadMiniDump*). These functions are configured to allow an  
9 arbitrary application to write and read a user minidump file. Experiments suggest  
10 that a user minidump file of approximately 100KB would be able to walk the  
11 callstack of all threads, read any active local variables, view parameters to  
12 functions, and view source code for the faulting application. Moreover, existing  
13 user-mode debugging tools, such as, Visual Studio and the like, can be easily  
14 debug a user minidump file. Furthermore, since the user minidump file format is  
15 completely specified and public, various developers can write custom applications  
16 to process the user minidump file as necessary.

17 Fig. 4 is a table 300 listing experimental results associated with user  
18 minidump process 200 as applied to certain applications during a failure. Here,  
19 user minidump files were created for Microsoft® Excel 2000, Outlook® 2000 and  
20 IE 5.0, using an x86 PC. As can be seen, the size of the resulting user minidump  
21 files for each of these applications is fairly small, ranging from 38 KB to 61 KB.  
22 Consequently, such user minidump files can be quickly saved and easily  
23 transported.

24 As shown in greater detail below, in certain implementations two system  
25 API calls have been added to DBGHELP of Windows® 2000, one to write a user

1 minidump file and one to read a user minidump file. Any application, including  
2 postmortem tools, debuggers or the application itself will be able to write a user  
3 minidump file at any time. In this manner, an application is explicitly allowed to  
4 write a user minidump file even when there is not an exception or failure.

5 The *WriteMiniDump* API is:

```
6
7     BOOL
8     WINAPI
9     WriteMiniDump(
10        IN HANDLE hProcess,
11        IN DWORD ProcessId,
12        IN HANDLE hFile,
13        IN MINIDUMP_TYPE DumpType,
14        IN PMINIDUMP_EXCEPTION_INFORMATION ExceptionParam,
15        IN PMINIDUMP_USER_STREAM_INFORMATION UserStream,
16        IN PMINIDUMP_CALLBACK_INFORMATION CallbackParam
17        );
18
```

19 The *WriteMiniDump* API writes a user-mode minidump to the file specified  
20 by *hFile*. The parameter *hProcess* is a process handle with full read and query  
21 permissions. *ProcessId* is a process identifier for the process. The *hFile* parameter  
22 is a handle to a file where the user minidump information is written. *DumpType* is  
23 used to identify the type of dump to generate. In other words, different types of  
24 minidumps supporting different feature sets can be supported. The *ExceptionParam* is  
25 an optional pointer to a structure describing the exception in the client that caused the mini-dump to be generated. If this parameter is NULL, no exception information will be written to the minidump file. The *UserStream* is an optional pointer to an array of *USER\_DATA\_ENTRY* structures. A user data entry is optional user information that is written to the dump file. The content of the user information is arbitrary and is never interpreted by the *WriteMiniDump*

1 routine. *UserStream* should be NULL if there is no optional user information.  
2 *CallbackParam* is an optional pointer to a user minidump callback data pointer  
3 and routine that *WriteMiniDump* can callback with extended information about the  
4 dump. The callback allows the caller of *WriteMiniDump* to completely control the  
5 granularity of information written to a user minidump file. If this pointer is NULL,  
6 no callbacks are performed.

7 The return values of the *WriteMiniDump* API are TRUE (1) upon success,  
8 and FALSE (0) upon failure.

9 The *ReadMiniDump* API is:

```
10
11     BOOL
12     WINAPI
13     ReadMinidumpStream (
14         IN PVOID BaseOfDump,
15         IN ULONG StreamNumber,
16         OUT PMINIDUMP_DIRECTORY * Dir, OPTIONAL
17         OUT PVOID * StreamPointer, OPTIONAL
18         OUT ULONG * StreamSize OPTIONAL
19     );
```

20 The *ReadMiniDump* API reads a field from a user-mode minidump. Since  
21 the user minidump file specification is well defined (e.g., see Fig. 5 below) the  
22 *ReadMiniDump* API can be considered strictly a convenience.

23 Here, *BaseOfDump* is the header of the mapped user minidump file,  
24 *StreamNumber* is the type of data to read from the mini-dump file, and *Dir*, on  
25 return, contains a pointer to the user minidump directory of type *DataType* file  
format. *StreamPointer*, on return, contains a pointer to the beginning of the mini-  
dump stream of type *DataType*. The type of data that this points to depends on the  
type of data requested in the *DataType* parameter. *StreamSize*, on return, contains  
the size of the stream.

1        The return values of the *ReadMiniDump* API are simply; TRUE (1) upon  
2        success, and FALSE (0) upon failure.

3        With this exemplary implementation in mind, a format for a user minidump  
4        file will now be described with reference to Fig. 5. The user minidump file 400  
5        includes a header section 402, a directory section 404 and a data section 406.  
6        These sections are indexed using RELATIVE VIRTUAL ADDRESSES (RVAs),  
7        which describe the location of a data member within a file. An RVA is an offset  
8        from the beginning of a file. In Fig. 5, RVA 408 points to the first directory in  
9        section 404, and RVAs 410(a-c) point to respective data 412(a-c) in section 406.

10       Since it is important for a user minidump file from one computer  
11       architecture to be readable on a computer of a different architecture, all structures  
12       in user minidump file 400 are therefore of the same size, regardless of the  
13       computer system they are generated on.

14       As shown in Fig. 5, the user minidump file format is similar to the Portable  
15       Executable image file format in that it specifies a set of directories that then point  
16       to the data (e.g., see *Microsoft Portable Executable Common Object File Format*  
17       *Specification Revision 6.0*: Visual C++ Business Unit, Microsoft Corporation,  
18       February 1999). Each directory entry in section 404 specifies the data type, the  
19       size of the data and an RVA 410 to where in the user minidump file 400 the data is  
20       located. In this example, only one directory of a specific type is allowed per file.

21       Those skilled in the art will recognize that the above described methods and  
22       arrangements are flexible in that a variety of operating systems and applications  
23       can be supported. The amount of resulting kernel minidump information or user  
24       minidump information can be increased or decreased as needed. For example, in  
25       certain implementations, the user minidump file can be reduced even further to

1 only include the faulting thread, faulting thread's callstack and context, module  
2 list, and reason for the crash. This can further reduce the user minidump file to  
3 several dozen bytes.

4 Additionally, in certain implementations, information can be added to the  
5 user minidump file to make it more usable by debuggers. One limitation of the  
6 user minidump process as described above is that it does not store any global  
7 variables or dynamically allocated memory. If needed, however, global variables  
8 could be saved by adding global data from the faulting process or, more  
9 sophisticatedly, by adding only selected regions of global data. These selected  
10 regions of global data can be selected by a custom-written application that has  
11 knowledge about where certain data is stored in the specific application. Also,  
12 selected data pages from the process could also be stored as needed to view active  
13 heap variables. The user minidump file format of Fig. 5 has the built-in flexibility  
14 to add these data while maintaining compatibility with previous versions.

15 Although some preferred embodiments of the various methods and  
16 arrangements of the present invention have been illustrated in the accompanying  
17 Drawings and described in the foregoing Detailed Description, it will be  
18 understood that the invention is not limited to the exemplary embodiments  
19 disclosed, but is capable of numerous rearrangements, modifications and  
20 substitutions without departing from the spirit of the invention as set forth and  
21 defined by the following claims.

1            **CLAIMS**

2

3        1. A method comprising:

4            a. determining when to generate a dump file; and

5            b. generating a dump file by gathering at least:

6                i. thread information for at least one running thread,

7                ii. context information for the thread,

8                iii. callstack information for the thread,

9                iv. process information for the process in which the thread is

10              running, and

11                v. information identifying a reason for generating the dump file.

12

13        2. The method as recited in Claim 1, wherein generating the dump file

14              further includes storing the dump file to a storage medium.

15

16        3. The method as recited in Claim 1, wherein generating the dump file

17              further includes gathering processor information about at least one

18              processor.

19

20

21        4. The method as recited in Claim 1, wherein determining when to

22              generate the dump file further includes determining that an exception

23              has occurred.

24

25

1       5. The method as recited in Claim 4, wherein the dump file does not  
2            further include any significant portion of a dynamically allocated  
3            memory.

4

5

6       6. The method as recited in Claim 5 wherein the dump file does not  
7            include any portion of a global initialized or uninitialized memory.

8

9       7. The method as recited Claim 5 wherein the dump file does not include  
10           any portion of the executable instructions used by the processor to  
11           execute the program.

12

13       8. The method as recited in Claim 1, wherein the dump file is a kernel  
14           minidump file associated with an operating system and the at least one  
15           running thread is the single thread which encountered the exception.

16

17

18       9. The method as recited in Claim 1, wherein the callstack information is  
19           a kernel stack.

20

21       10. The method as recited in Claim 1, wherein the process information  
22           identifies the process that initiated the thread.

23

24       11. The method as recited in Claim 1, further comprising:

25

- a. allocating a buffer space in memory during an initialization process;  
and

b. reserving space on a storage medium drive suitable for writing the contents of the buffer.

12. The method as recited in Claim 11, wherein:

- a. generating the dump file further includes initially storing the thread information, the context information, the callstack information, the process information, and the information identifying the reason for generating the dump file to the buffer space, and then copying the dump file from the buffer space to the storage medium as a dump file; and

b. upon system re-initialization, transferring the dump file from the storage medium to at least one external computer.

13. The method as recited in Claim 12, further comprising upon re-initialization, after having stored the dump file to the storage medium, accessing the dump file on the storage medium and using at least a portion of the dump file to further understand an exception that was at least one reason for generating the dump file.

14. The method as recited in Claim 1, wherein the dump file is a user minidump file associated with at least one non-operating system program.

1

2 15. The method as recited in Claim 1, wherein generating the dump file  
3 further includes gathering callstack information for all running threads.

4

5 16. The method as recited in Claim 1, wherein the callstack information is  
6 a user callstack.

7

8 17. The method as recited in Claim 1, wherein generating the dump file  
9 further includes gathering processor context information for all running  
10 threads.

11

12 18. The method as recited in Claim 1, wherein generating the dump file  
13 further includes gathering a listing of all loaded modules for the  
14 faulting application program.

15

16 19. The method as recited in Claim 1, wherein the dump file is a directory  
17 indexed file that uses relative virtual addresses (RVAs).

18

19 20. A computer-readable medium having computer-executable instructions  
20 for performing steps comprising:  
21  
a. determining when to generate a dump file; and  
22  
b. generating a dump file by gathering at least:  
23  
i. thread information for at least one running thread,  
24  
ii. context information for the thread,  
25  
iii. callstack information for the thread,

- iv. process information for the process in which the thread is running, and
- v. information identifying a reason for generating the dump file.

21. The computer-readable medium as recited in Claim 20, wherein generating the dump file further includes storing the dump file to a storage medium.

22. The computer-readable medium as recited in Claim 20, wherein generating the dump file further includes gathering processor information about at least one processor.

23. The computer-readable medium as recited in Claim 20, wherein determining when to generate the dump file further includes determining that an exception has occurred.

24. The computer-readable medium as recited in Claim 23, wherein the dump file does not further include any significant portion of a dynamically allocated memory.

25. The computer-readable medium as recited in Claim 24 wherein the dump file does not include any portion of a global initialized or uninitialized memory.

1       26. The computer-readable medium as recited Claim 24 wherein the dump  
2       file does not include any portion of the executable instructions used by  
3       the processor to execute the program.

4

5       27. The computer-readable medium as recited in Claim 20, wherein the  
6       dump file is a kernel minidump file associated with an operating  
7       system and the at least one running thread is the single thread which  
8       encountered the exception.

9

10      28. The computer-readable medium as recited in Claim 20, wherein the  
11       callstack information is a kernel stack.

12

13      29. The computer-readable medium as recited in Claim 20, wherein the  
14       process information identifies the process that initiated the thread.

15

16      30. The computer-readable medium as recited in Claim 20, further  
17       comprising computer-executable instructions for performing steps of:  
18       allocating a buffer space in memory during an initialization process;  
19       and  
20       reserving space on a storage medium drive suitable for writing the  
21       contents of the buffer.

22

23      31. The computer-readable medium as recited in Claim 30, wherein:  
24       generating the dump file further includes initially storing the thread  
25       information, the context information, the callstack information, the

1 process information, and the information identifying the reason for  
2 generating the dump file to the buffer space, and then copying the dump  
3 file from the buffer space to the storage medium as a dump file; and  
4 upon system re-initialization, transferring the dump file from the storage  
5 medium to at least one external different computer.

6

7

8 32. The computer-readable medium as recited in Claim 31, further  
9 comprising computer-executable instructions for performing steps of,  
10 upon re-initialization after having stored the dump file to the storage  
11 medium, accessing the dump file on the storage medium and using at  
12 least a portion of the dump file to further understand an exception that  
13 was at least one reason for generating the dump file.

14

15 33. The computer-readable medium as recited in Claim 20, wherein the  
16 dump file is a user minidump file associated with at least one non-  
17 operating system program.

18

19 34. The computer-readable medium as recited in Claim 20, wherein generating  
20 the dump file further includes gathering callstack  
21 information for all running threads.

22

23 35. The computer-readable medium as recited in Claim 20, wherein the  
24 callstack information is a user callstack.

25

1       36. The computer-readable medium as recited in Claim 20, wherein  
2       generating the dump file further includes gathering processor context  
3       information for all running threads.

4

5       37. The computer-readable medium as recited in Claim 20, wherein  
6       generating the dump file further includes gathering a listing of all  
7       loaded modules for the faulting application program.

8

9       38. The computer-readable medium as recited in Claim 20, wherein the  
10      dump file is a directory indexed file that uses relative virtual addresses  
11      (RVAs).

12

13

14       39. An arrangement comprising memory, a data storage drive configured to  
15      write data files to at least one data storage medium, and a processor  
16      operatively coupled to the memory and the data storage drive and  
17      configured to:  
18           a. determine when to generate a dump file; and  
19           b. generate a dump file by gathering at least:  
20               i. thread information for at least one running thread,  
21               ii. context information for the thread,  
22               iii. callstack information for the thread,  
23               iv. process information for the process in which the thread is  
24              running, and information identifying a reason for generating  
25              the dump file.

1

2 40. The arrangement as recited in Claim 39, wherein generating the dump

3 file further includes storing the dump file to a storage medium.

4

5 41. The arrangement as recited in Claim 39, wherein generating the dump

6 file further includes gathering processor information about at least one

7 processor.

8

9 42. The arrangement as recited in Claim 39, wherein determining when to

10 generate the dump file further includes determining that an exception

11 has occurred.

12

13 43. The arrangement as recited in Claim 43, wherein the dump file does

14 not further include any significant portion of a dynamically allocated

15 memory.

16

17 44. The arrangement as recited in Claim 43 wherein the dump file does not

18 include any portion of a global initialized or uninitialized memory.

19

20 45. The arrangement as recited Claim 43 wherein the dump file does not

21 include any portion of the executable instructions used by the processor

22 to execute the program.

23

24 46. The arrangement as recited in Claim 39, wherein the dump file is a

25 kernel minidump file associated with an operating system and the at

least one running thread is the single thread which encountered the exception.

47. The arrangement as recited in Claim 39, wherein the callstack information is a kernel stack.
48. The arrangement as recited in Claim 39, wherein the process information identifies the process that initiated the thread.
49. The arrangement as recited in Claim 39, further comprising computer-executable instructions for performing steps of:
  - allocating a buffer space in memory during an initialization process;
  - and
  - reserving space on a storage medium drive suitable for writing the contents of the buffer.
50. The arrangement as recited in Claim 49, wherein:
  - generating the dump file further includes initially storing the thread information, the context information, the callstack information, the process information, and the information identifying the reason for generating the dump file to the buffer space, and then copying the dump file from the buffer space to the storage medium as a dump file; and
  - upon system re-initialization, transferring the dump file from the storage medium to at least one external computer.

1        51. The arrangement as recited in Claim 50, further comprising computer-  
2        executable instructions for performing steps of, upon re-initialization  
3        after having stored the dump file to the storage medium, accessing the  
4        dump file on the storage medium and using at least a portion of the  
5        dump file to further understand an exception that was at least one  
6        reason for generating the dump file.

7

8        52. The arrangement as recited in Claim 39, wherein the dump file is a user  
9        minidump file associated with at least one non-operating system  
10       program.

11

12       53. The arrangement as recited in Claim 39, wherein generating the dump  
13       file further includes gathering callstack information for all running  
14       threads.

15

16       54. The arrangement as recited in Claim 39, wherein the callstack  
17       information is a user callstack.

18

19       55. The arrangement as recited in Claim 39, wherein generating the dump  
20       file further includes gathering processor context information for all  
21       running threads.

22

23       56. The arrangement as recited in Claim 39, wherein generating the dump  
24       file further includes gathering a listing of all loaded modules for the  
25       faulting application program.

1  
2 57. The arrangement as recited in Claim 39, wherein the dump file is a  
3 directory indexed file that uses relative virtual addresses (RVAs).

4  
5 58. A method for generating a minimal dump file, the method comprising:  
6     a. detecting an exception; and  
7     b. outputting:  
8         i. information on a faulting thread and an associated process,  
9             and  
10         ii. a list of loaded modules.

11  
12 59. The method as recited in Claim 58, further comprising storing the  
13 minimal dump file to a storage medium.

14  
15 60. The method as recited in Claim 58, further comprising transporting the  
16 minimal dump file using a communication resource.

17  
18 61. A method of communicating between a client process and a server  
19 process in a distributed processing system, comprising:  
20     a. issuing, by the client process, a write dump file call having a  
21         plurality of call parameters comprising a process handle, a  
22         process identifier, a handle to a file where dump file information  
23         is to be written, and a dump type identifier;  
24     b. receiving, by the server process, the write dump file call and  
25         parsing the call to retrieve the parameters; and

c. issuing, by the server process, a write dump file call acknowledgment providing a true-false indication.

62. The method as recited in Claim 61, wherein the plurality of call parameters further includes a pointer to a structure describing an exception in the client that caused the dump file to be generated.

63. The method as recited in Claim 61, wherein the plurality of call parameters further includes a pointer to an array of user data entry structures.

64. The method as recited in Claim 61, wherein the plurality of call parameters further includes a pointer to a dump file callback data pointer.

65. A method of communicating between a client process and a server process in a distributed processing system, comprising:

- a. issuing, by the client process, a read dump file call having a plurality of call parameters comprising a header of a dump file and a data type identifier of data to read from a dump file;
- b. receiving, by the server process, the read dump file call and parsing the call to retrieve the parameters; and
- c. issuing, by the server process, a read dump file call acknowledgment providing a true-false indication and a plurality of call return parameters comprising a pointer to a beginning of a

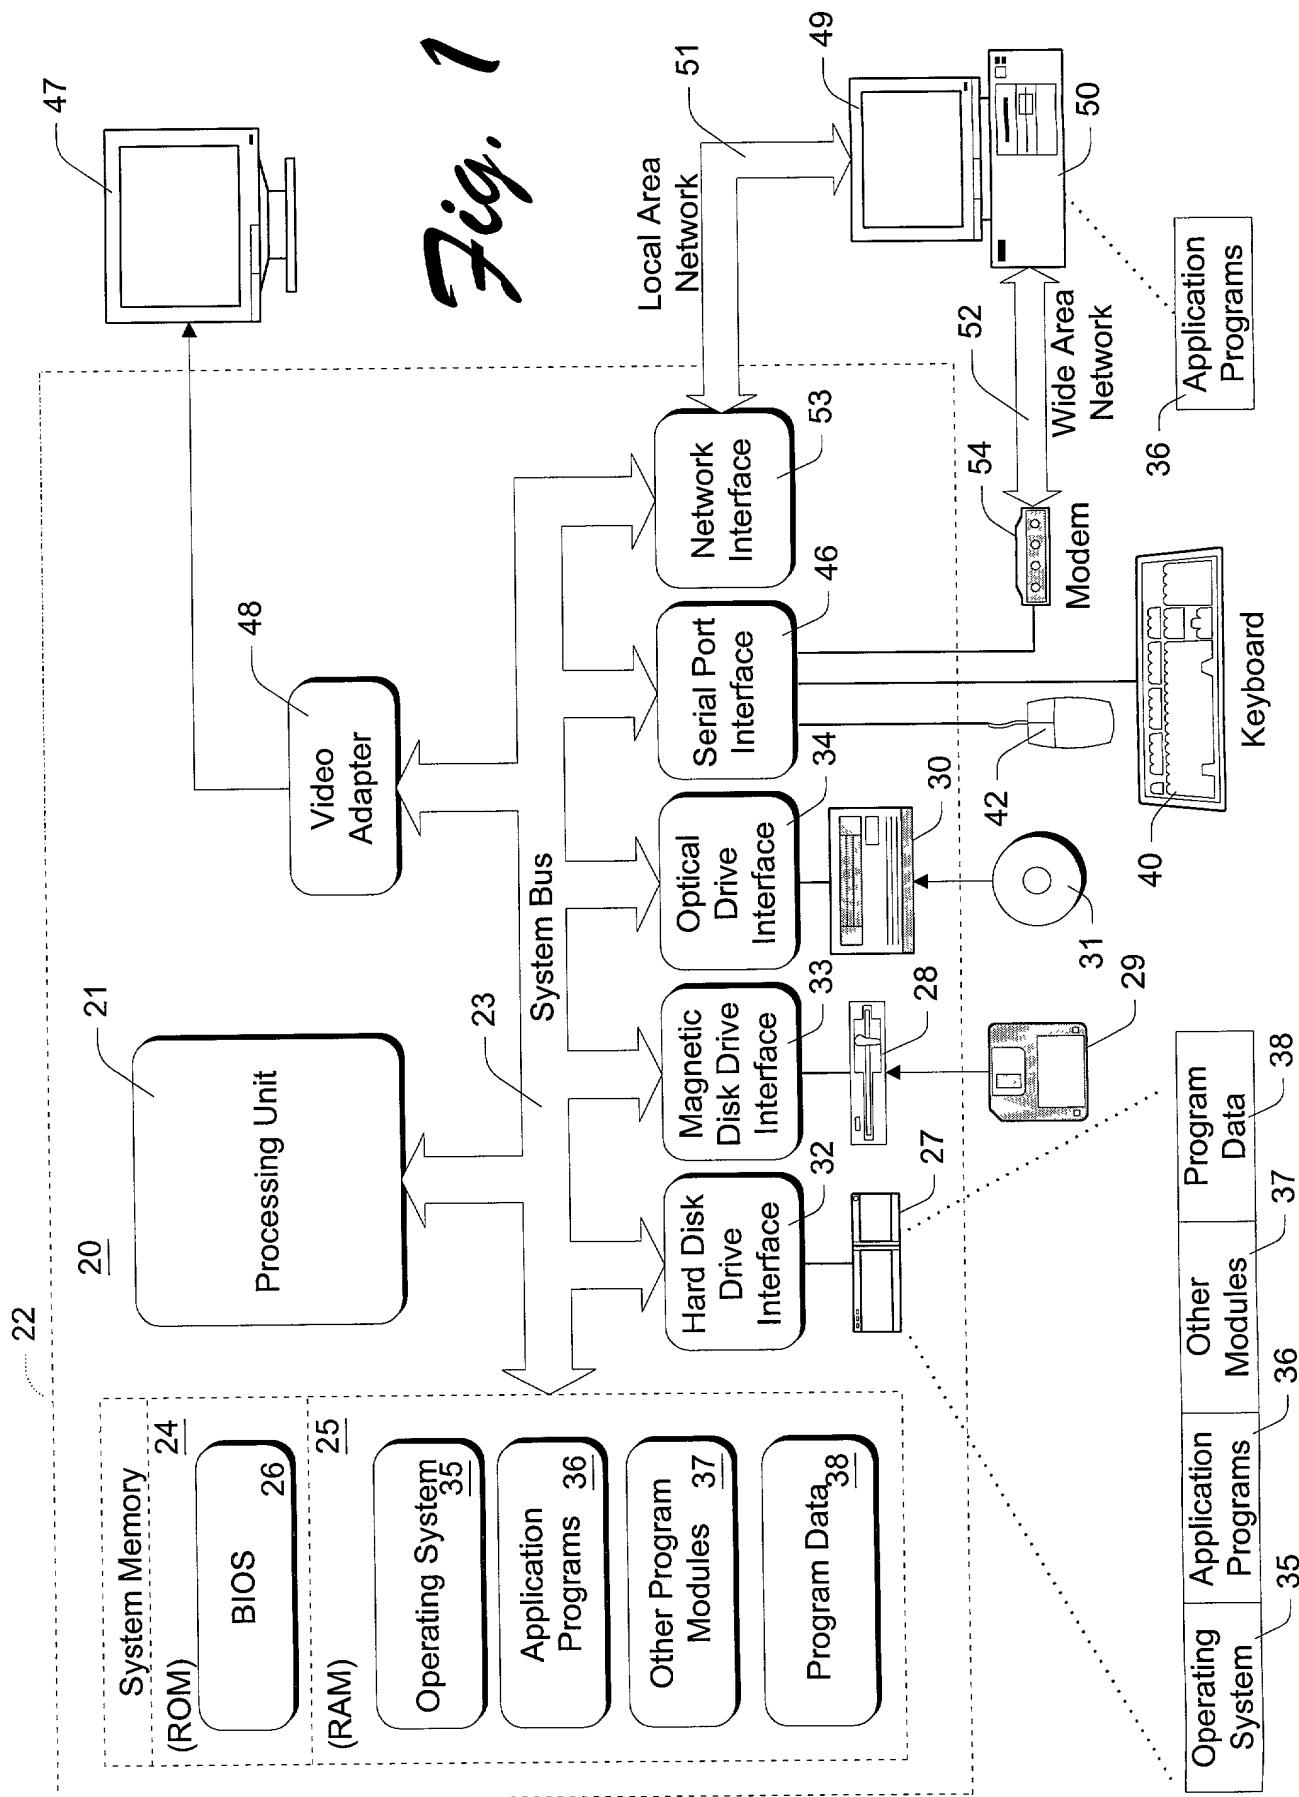
1 dump stream, and a stream size identifying the size of the dump  
2 stream.

3

4 66. The method as recited in Claim 65, wherein the plurality of call return  
5 parameters further includes a pointer to a dump file directory.

1      **ABSTRACT**

2      Methods and arrangements are provided that substantially reduce the  
3      requisite amount of data required to conduct postmortem analysis following a  
4      computer failure. The methods and arrangements can be advantageously  
5      configured to allow for rapid online user support for a variety of users, computing  
6      devices, operating systems, applications, and the like. One method includes  
7      determining when to generate a dump file, and generating a dump file by gathering  
8      thread, callstack and thread context information for the running thread, process  
9      identifying information associated with the running thread, and information  
10     identifying the reason for generating the dump file. The resulting dump file is  
11     then stored to a storage medium and accessed during subsequent analysis. The  
12     dump file can be a kernel minidump file that is associated with an operating  
13     system program failure, in which case the running thread is the single thread that  
14     was running when the failure occurred. The kernel minidump file would include  
15     the kernel callstack and the process identifying information that would identify the  
16     process that initiated the single thread. The method is further applicable to non-  
17     operating system programs, wherein a user minidump file is generated by also  
18     gathering callstack information for all running threads, thread context information  
19     for all running threads, and a listing of all loaded modules for the faulting non-  
20     operating system program.



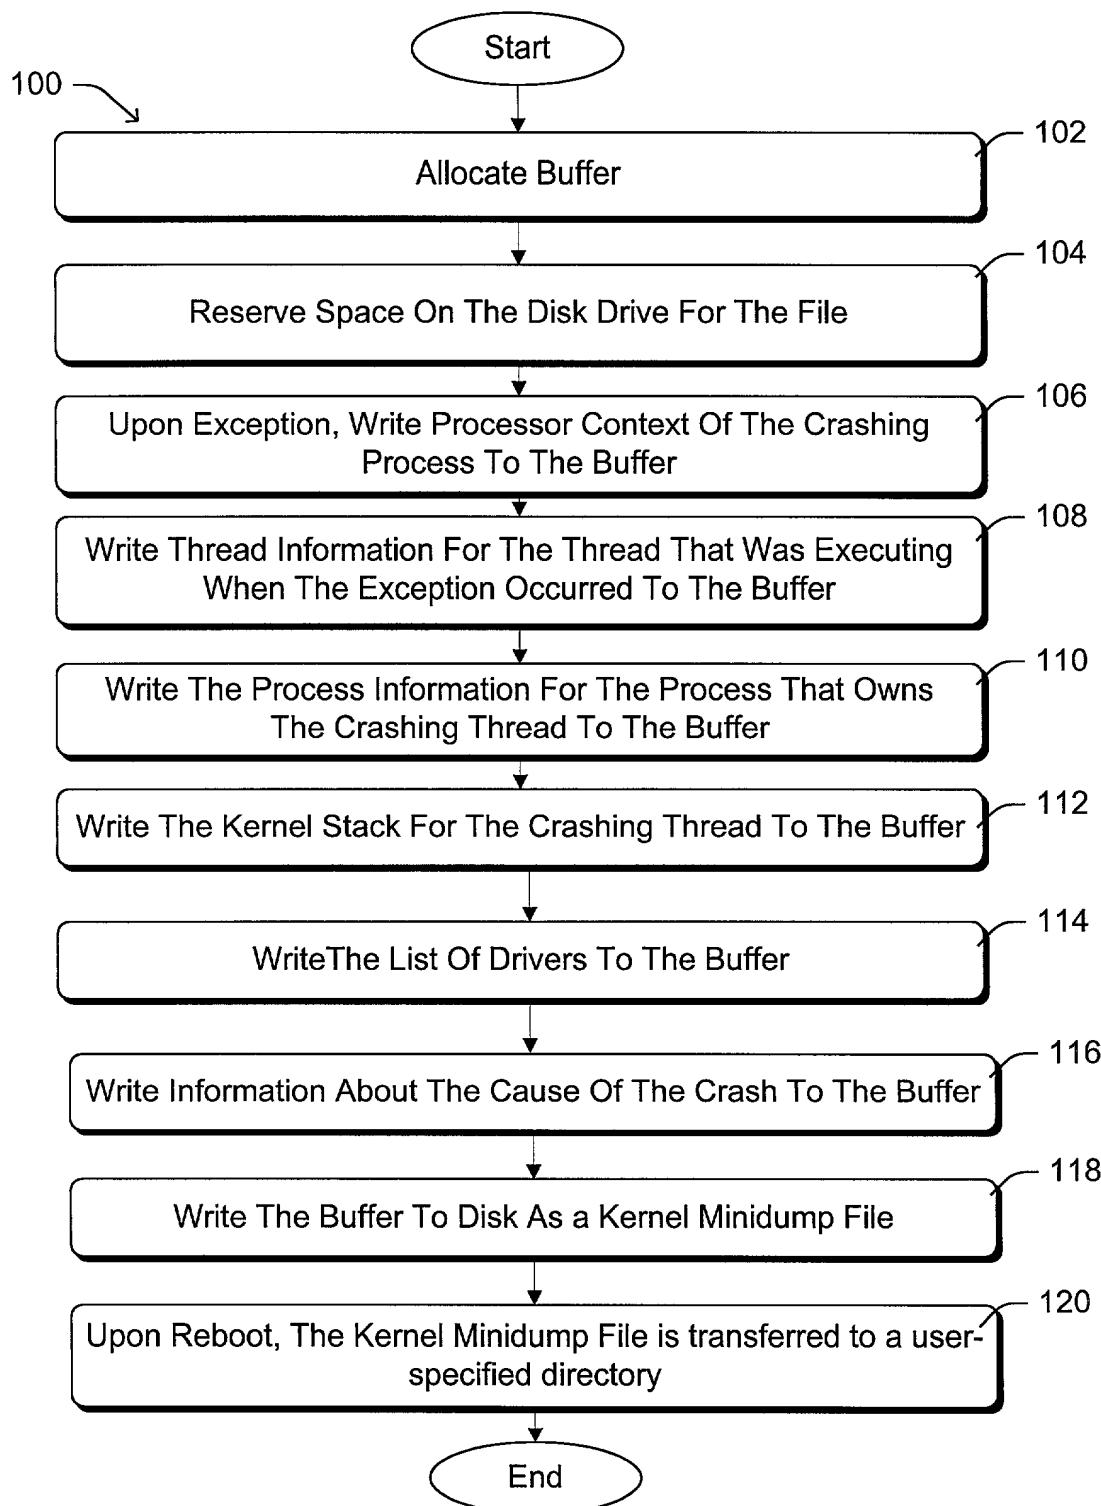


Fig. 2

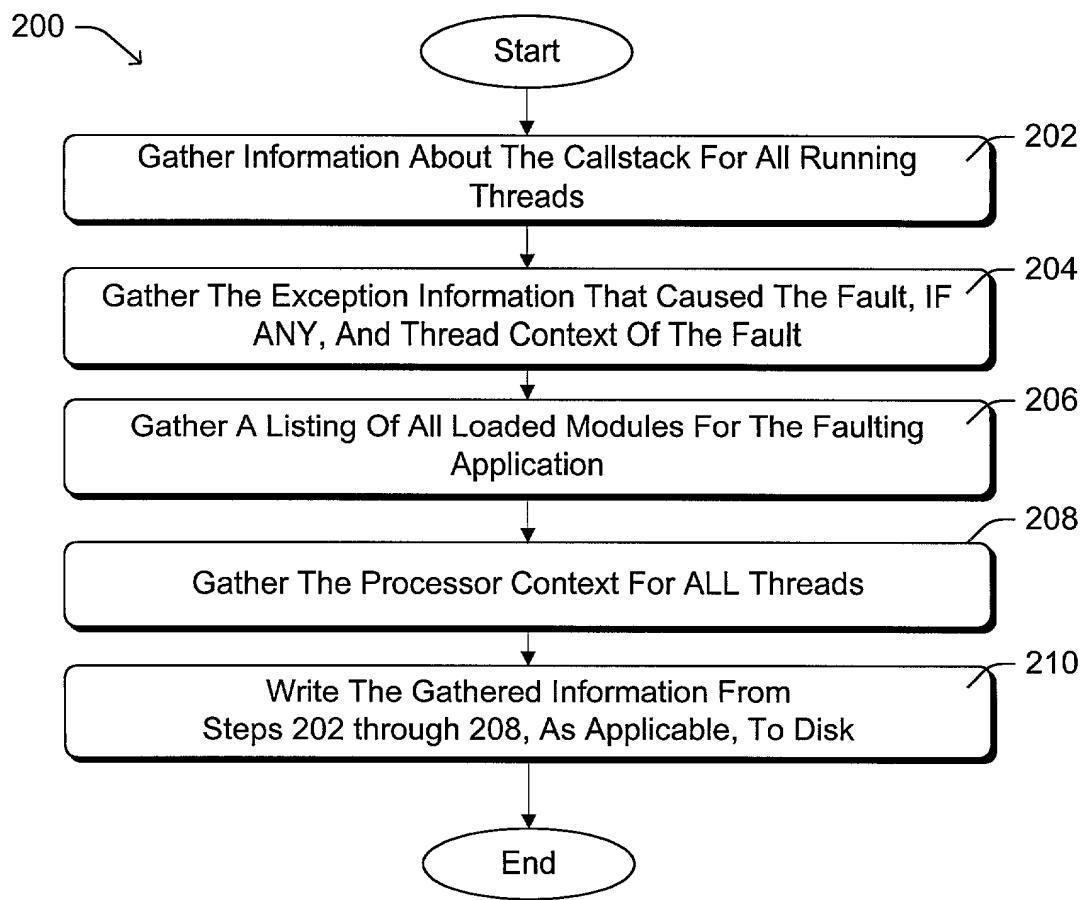


Fig. 3

300 ↘

	Excel	Outlook	IE
Number of threads	10	14	12
Largest Stack size in KB	2	2	8
Total Stack Size in KB	5	5	17
Number of Modules	56	79	82
Size of Mini-dump File in KB	38	50	61

Fig. 4

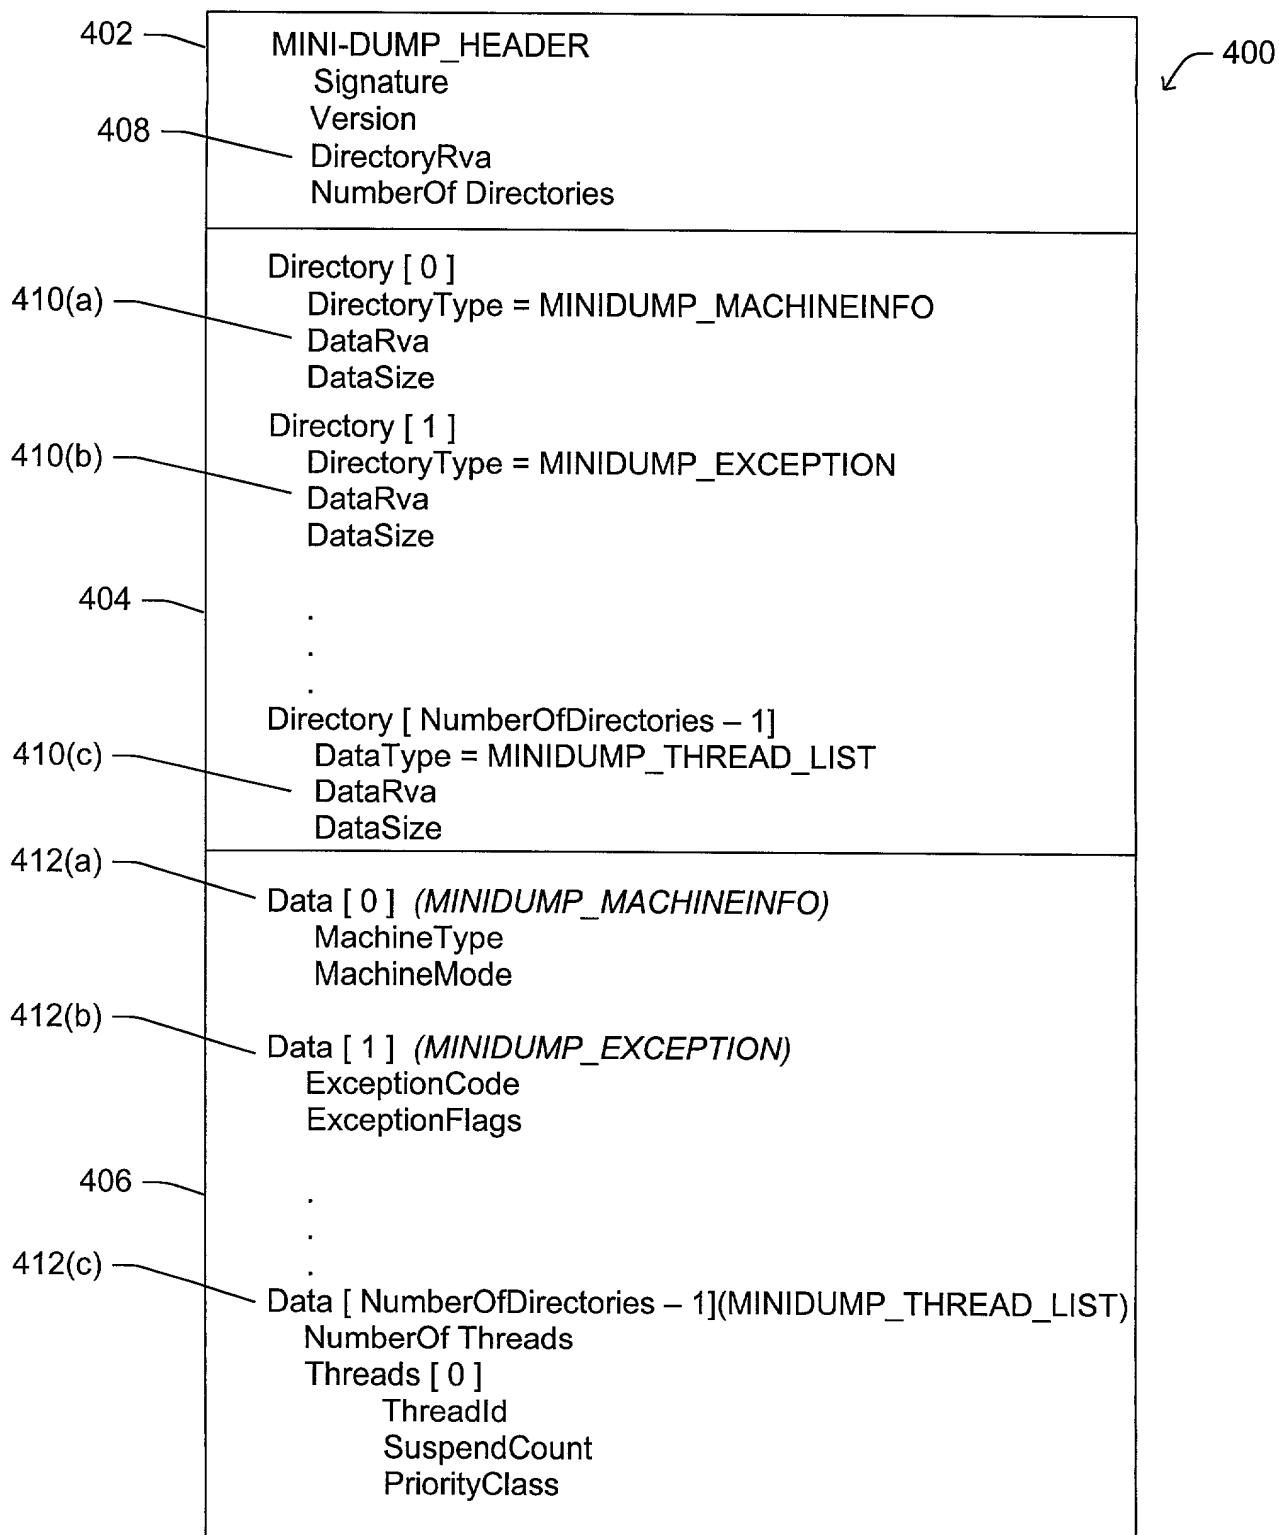


Fig. 5

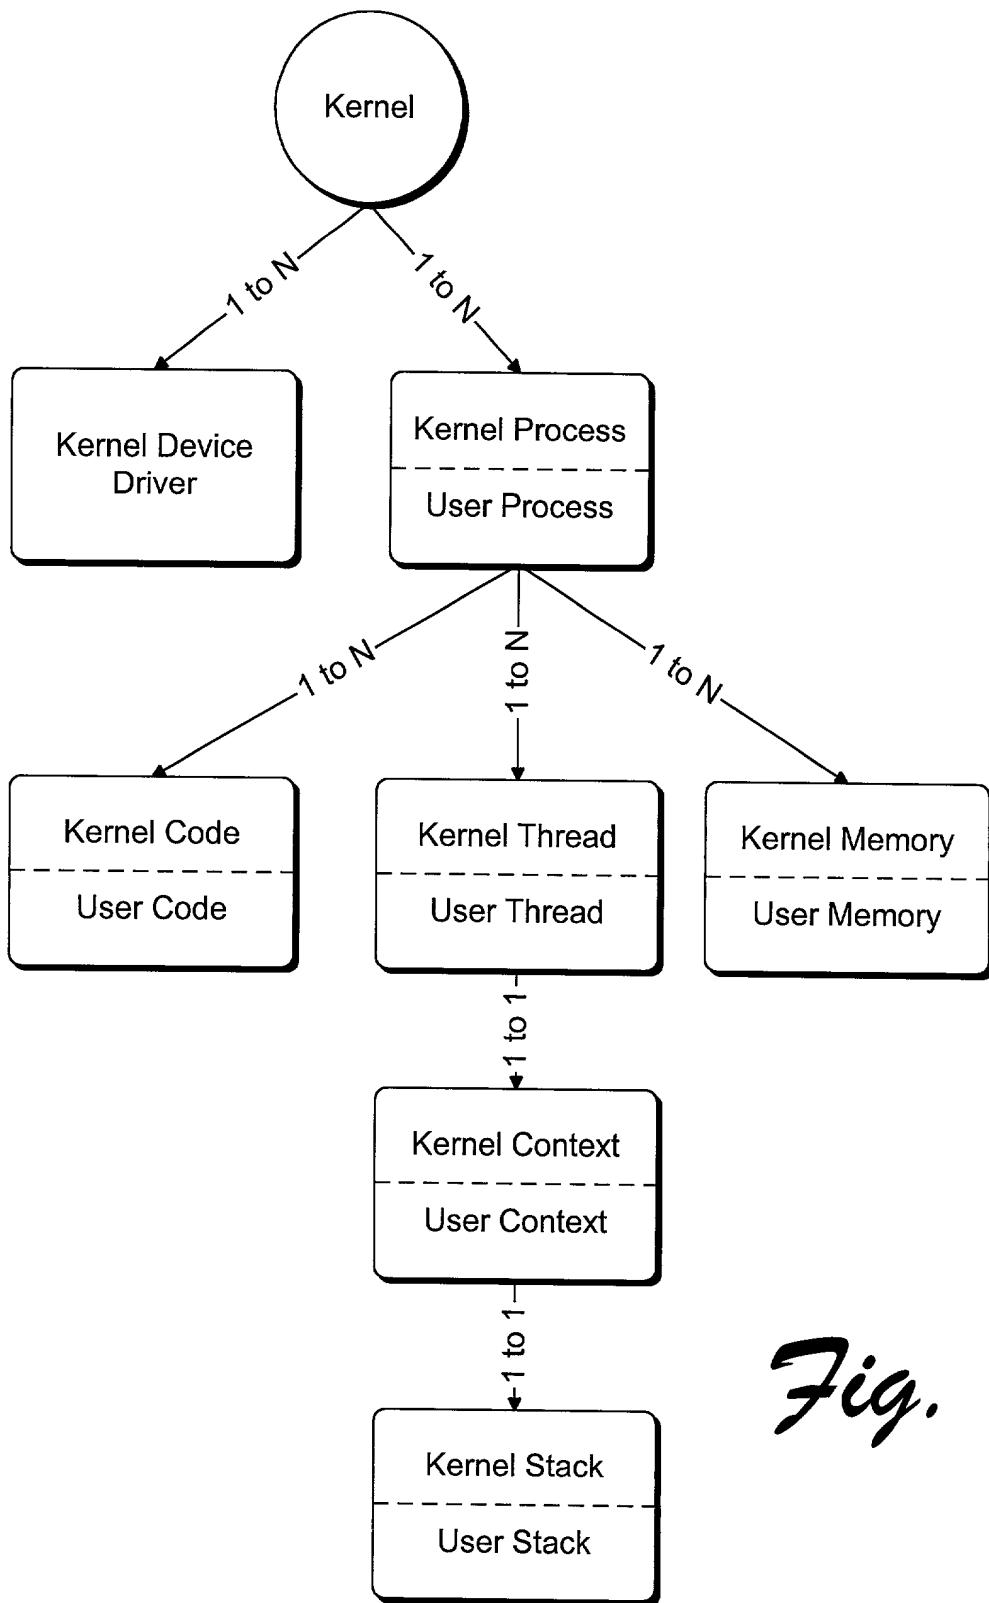


Fig. 6